



US005878238A

**United States Patent** [19]  
**Gan et al.**

[11] **Patent Number:** **5,878,238**  
 [45] **Date of Patent:** **Mar. 2, 1999**

[54] **TECHNIQUE FOR SUPPORTING SEMI-COMPLIANT PCI DEVICES BEHIND A PCI-TO-PCI BRIDGE**

[75] **Inventors:** Doron Gan; Jeff Savage, both of Austin, Tex.

[73] **Assignee:** Dell USA, L.P., Round Rock, Tex.

[21] **Appl. No.:** 908,650

[22] **Filed:** Aug. 7, 1997

#### Related U.S. Application Data

[62] Division of Ser. No. 590,461, Jan. 23, 1996.

[51] **Int. Cl.<sup>6</sup>** ..... G06F 1/30; G06F 13/40

[52] **U.S. Cl.** ..... 395/308; 395/309; 395/828; 395/183.12

[58] **Field of Search** ..... 395/280-284, 395/306-309, 828-834, 183.12

#### [56] References Cited

##### U.S. PATENT DOCUMENTS

5,018,146	5/1991	Sexton	371/37.1
5,325,369	6/1994	Kim	371/29.1
5,375,211	12/1994	Maruyama et al.	395/325
5,379,384	1/1995	Solomon	395/325
5,483,468	1/1996	Chen et al.	364/551.01
5,555,425	9/1996	Zeller et al.	
5,560,016	9/1996	Fiebrich et al.	

5,568,619	10/1996	Blackledge et al.	395/281
5,581,714	12/1996	Amini et al.	395/308
5,594,873	1/1997	Garrett	
5,608,876	3/1997	Cohen et al.	395/281
5,611,045	3/1997	Basile	395/184.01
5,615,323	3/1997	Engel et al.	395/140
5,619,728	4/1997	Jones et al.	
5,632,021	5/1997	Jennings et al.	395/309
5,636,342	6/1997	Jeffries	
5,644,470	7/1997	Benedict et al.	361/686
5,689,726	11/1997	Lin	395/830
5,692,219	11/1997	Chan et al.	

*Primary Examiner*—Ayaz R. Sheikh


*Assistant Examiner*—Jigar Pancholi



*Attorney, Agent, or Firm*—Haynes and Boone, LLP




#### [57] ABSTRACT

Method and apparatus for detecting the presence of a semi-compliant PCI device in a secondary expansion slot of a PC and instructing the user to reinsert the device into one of the primary slots are disclosed. In one embodiment, upon detection of a semi-compliant PCI device in a secondary slot, a video image instructing the user to reinsert the device into one of the primary slots is displayed on a display of the PC. Operation remains suspended until the device is relocated to a primary slot. In a presently preferred embodiment, a hardware enhancement to a PCI-to-PCI bridge connecting a primary PCI bus to a secondary PCI bus enables the device to operate flawlessly on the secondary PCI bus, such that the user remains unaware of the otherwise undesirable situation.

24 Claims, 2 Drawing Sheets







# TECHNIQUE FOR SUPPORTING SEMI-COMPLIANT PCI DEVICES BEHIND A PCI-TO-PCI BRIDGE

This is a divisional of copending application Ser. No. 08/590,461 filed on Jan. 23, 1996.

## TECHNICAL FIELD

The invention relates generally to semi-compliant PCI devices and, more particularly, to a technique for supporting such devices behind a PCI-to-PCI bridge.

## BACKGROUND OF THE INVENTION

It is well known that the number of devices that can be supported on a single peripheral component interconnect, or "PCI," bus is limited. In particular, the number of device loads that can be connected to one such bus is, under normal circumstances, approximately ten. Typically, one load is allocated for a motherboard device and two loads each are allocated for expansion slots. In conventional PC systems, three of the ten loads are typically used for three motherboard-based devices, respectively, one load is allocated to a riser connector, when a riser card implementation is used to support expansion slots is used, leaving six loads for three expansion slots into which expansion cards can be plugged. Accordingly, to enable more than two or three PCI expansion slots to be supported on a single PC, it has been known include in a PC an additional PCI bus (the "secondary PCI bus"), which is connected to the first PCI bus (the "primary PCI bus") via a device commonly designated as a PCI-to-PCI bridge.

A PCI-to-PCI bridge is typically implemented as an integrated circuit ("IC") chip connected to a primary PCI bus in front of the bridge and forming a secondary PCI bus behind the bridge. The bridge only takes up one device load on the primary bus and provides for an additional ten device loads behind the bridge via the secondary bus. These additional loads are almost totally electrically isolated from the primary bus.

FIG. 1 illustrates a PC 10 comprising a typical PCI configuration wherein a PCI-to-PCI bridge 12 is used to increase the number of PCI devices that can be supported. In particular, the PC 10 includes a host 14, which will include a host CPU, system memory, and ROM BIOS, residing on a host bus 16. The host bus 16 is connected to a primary PCI bus 18 via a host-to-PCI bridge 20. The primary PCI bus 18 is further connected to a secondary PCI bus 22, on which a display 23 resides, via the PCI-to-PCI bridge 12. In the illustrated embodiment, a single PCI device 26 resides on the primary PCI bus 18. In addition, two primary expansion slots 28a, 28b, for enabling additional PCI devices to be connected to the primary PCI bus 18, are provided thereon. Similarly, two PCI devices 32, 34, reside directly on the secondary PCI bus 22 and three secondary expansion slots 36a-36c are provided on the secondary PCI bus 22.

It should be understood that more than the number of expansion slots shown in FIG. 1 may be connected to the buses 18 and 22. Moreover, although not shown, will be recognized that the slots 28a-28b, 36a-36c, may be connected to the respective bus 18, 22, via an appropriate riser card.

In operation, each PCI 2.x-compliant device, such as the devices 26, 32 and 34, residing on a PCI bus requests a certain address range through which other devices can access the PCI device. Each device requests some number of consecutive addresses and the host 14 assigns chunks of the

I/O space to the device. For example, if the device 26 requests a 256 byte address space, the host 14 may assign the device the contiguous address space beginning at address FECh. The device 26 will then use the assigned address as a base, or lower, address limit and the assigned address plus 256 bytes as the upper address limit. Thereafter, any writes to or reads from this I/O space will be claimed by the device 26 and the device 26 will respond appropriately. A PCI-to-PCI bridge, such as the bridge 12, being a PCI device, operates in the same fashion. Specifically, it requests a contiguous address space, in particular, 4K, for its address space. The bridge 12 then divides this space and issues segments of it to the devices residing on the secondary PCI bus 22, such as devices 32, 34, and slots 36a-36c. In this manner, each device and slot that resides on the secondary bus 22 will be assigned its own address space comprising part of the address space assigned to the bridge 12.

When a device on the primary PCI bus 18, such as the device 26, wants to communicate with a device on the secondary PCI bus 22, such as the device 32, the device 26 will attempt to write to or read from the address space assigned to the bridge 12 and allocated by the bridge 12 to the device 32. The bridge 12 will accept the transaction and forward it to the device 32.

There are certain problems inherent in the use of a conventional PCI-to-PCI bridge in the manner described above. First, there are many PCI devices, in particular video cards and IDE cards, that are not completely PCI compliant. Such devices are hereinafter referred to as "semi-compliant PCI devices." Semi-compliant PCI devices have inherited an addressing scheme from the industry standard architecture ("ISA") standard that allows an expansion card to specify a noncontiguous address space. In addition, semi-compliant PCI devices have designated addresses that do not apply to the PCI remapping feature as defined by PCI 2.x specifications. Because PCI-to-PCI bridges that are currently commercially available only allocate to themselves a single contiguous I/O address range and cannot selectively accept certain single noncontiguous addresses, semi-compliant devices will not operate properly behind current PCI-to-PCI bridges. Moreover, although current PCI-to-PCI bridges do possess a special feature that enables them selectively to forward certain noncontiguous addresses in the VGA space according to the VGA specification, they lack two or more addresses that are not in the specification, but that most video devices require. As a result, the behavior observed by a user as a result of plugging a video card into one of the secondary expansion slots 36a-36c, which are generally externally indistinguishable from the primary expansion slots 28a, 28b, would be a lack of a video image; i.e., a blank screen. Clearly, this is an unacceptable result.

Additionally, some PCI devices, although capable of functioning behind a PCI-to-PCI bridge, incur a performance penalty in doing so. This is primarily due to the fact that any data must go through an intermediate device (i.e., the bridge 12) on the way to and from the device, thereby resulting in a latency. For ease of discussion, such devices will also be referred to herein as "semi-compliant PCI devices." Unless notified, a user might remain unaware that the device would operate more efficiently if inserted in one of the primary PCI slots.

Therefore, what is needed is a technique for detecting the insertion of a semi-compliant PCI device into a secondary expansion slot and instructing regarding same.

## SUMMARY OF THE INVENTION

The present invention, accordingly, provides a method and apparatus for detecting the presence of a semi-compliant

PCI device in a secondary expansion slot of a PC and instructing the user to reinsert the device into one of the primary slots, thereby overcoming or reducing disadvantages and limitations associated with prior methods and systems. In one embodiment, upon detection of a semi-compliant PCI device in a secondary slot, a video image instructing the user to reinsert the device into one of the primary slots is displayed on a display of the PC. Operation of the device will remain disabled until it is moved to a primary slot.

In other alternative embodiments, the user may be directed to move the semi-compliant PCI device from a secondary slot to a primary slot via a visual indicator, such as an illuminated LED, or an audio indicator, such as a series of beeps or instructions. As with the embodiment described above, operation of the device will remain disabled until it is moved to a primary slot.

In another, and presently preferred, embodiment, the invention comprises a hardware enhancement to a PCI-to-PCI bridge that allows fine grain programmable control of a set of I/O address ranges. In one embodiment, the enhancement circuit comprises four register sets, each set comprising a base and a limit register, for enabling a user to program several I/O ranges, and a configuration register for indicating a mode of operation of the circuit. When an address received by the PCI-to-PCI bridge lies within the I/O range defined by any one of the register sets, the action taken is determined by the mode indicated by the configuration register. In particular, if the configuration register indicates a first mode of operation, the I/O cycle is simply forwarded by the bridge. In contrast, if the configuration register indicates a second mode of operation, an interrupt is generated on an interrupt pin of the chip, allowing the BIOS to handle an exception. In this embodiment, it is possible for use of a semi-compliant PCI device to be enabled without requiring its relocation to a primary slot.

A technical advantage achieved with the invention is that it helps to insure that a user is made aware of the fact that a device that is inoperable behind a PCI-to-PCI bridge, such as most video cards, should be moved to a primary PCI expansion slot to enable use thereof.

Another technical advantage achieved with the invention is that it ensures that a user is made aware of the fact that certain devices, when plugged into a secondary PCI expansion slot do not function as efficiently as they would if plugged into a primary slot.

Yet another technical advantage achieved with the invention is that, in the preferred embodiment, it enables a video card to be used from a secondary PCI expansion slot.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a system block diagram of a PC comprising a conventional PCI bus configuration.

FIGS. 2A and 2B are flowcharts illustrating a first method of implementing the technique of the present invention for detecting the presence of a semi-compliant PCI device in a secondary expansion slot and responding to same.

FIG. 3 is a flowchart illustrating an alternative method of implementing the technique of the present invention.

FIG. 4 is a block diagram of a hardware configuration embodying a preferred implementation of the present invention.

#### DESCRIPTION OF THE PREFERRED EMBODIMENTS

As described above, FIG. 1 is a system block diagram of a PC 10 comprising a conventional PCI bus configuration.

Referring now to FIGS. 2A and 2B, a flowchart of a method of detecting the presence of a semi-compliant PCI device in a secondary expansion slot and informing the user of this condition via a video image or other indicator. In particular, the method illustrated in FIGS. 2A and 2B utilizes an exception handler to detect that a semi-compliant PCI device is plugged into one of the secondary slots 36a-36c and instruct the user to move the device from the secondary slot to one of the primary slots 28a, 28b. It will be appreciated that instructions for execution by the host processor for implementing the method illustrated in FIGS. 2A and 2B, as well as the method illustrated in FIG. 3, are stored in a memory device of the PC 10.

In step 200, during POST of the PC 10, a determination is made whether there is a non-compliant PCI device, specifically, a video card, on the secondary bus 22. If not, execution of the routine terminates in step 201 and POST proceeds as usual; otherwise, execution proceeds to step 202. In step 202, the host processor of the host 14 is set to trap on I/O writes to VGA address space. As a result of this step, any attempted write cycle to VGA address space will result in the execution of an exception handler, described in detail with reference to FIG. 2B. In step 204, an option, or BIOS, ROM (not shown) of the detected video card is executed and the card is initialized. In step 206, a message instructing the user to remove the video card from the secondary expansion slot and reinsert it into a primary expansion slot is displayed on the display 15. Operation remains suspended until the card is moved, it being understood that the PC 10 must be turned off prior to moving the card. Execution of the routine then terminates in step 201.

As previously indicated, as a result of step 202, any attempted write cycle to an address in the VGA address space will result in the execution of an exception handler, the operation of which will now be described with reference to FIG. 2B. In step 220, the option ROM of the video card is activated and the I/O address space thereof is read. In step 222, an address look-up table is generated. In step 224, a determination is made whether the address of the write cycle is part of the address space of the card as specified in the look-up table. If not, execution returns to step 220; otherwise, execution proceeds to step 226, in which the bridge 12 is remapped to handle the address. In step 228, the cycle is forwarded to the video card and the card is initialized and in step 230, the bridge 13 is remapped to its original mapping. Finally, in step 232, a return from exception is executed.

In this manner, the video card may be initialized and used to instruct the user, via a message displayed on the display 15, to move the video card to a primary expansion slot 28a, 28b, before further use of the PC 10.

As shown in FIG. 3, in an alternative embodiment, during POST, a determination is made in step 300 whether a semi-compliant PCI device is detected in a secondary expansion slot 36a-36c. This step may be performed, for example, by comparing the identity of cards in the secondary PCI slots 36a-36c with a list of semi-compliant PCI card stored in a memory device of the host 14. In this manner, any device that is not capable of functioning, or that is capable of functioning, but not optimally, behind the bridge 12, may be identified in the list and designated as "semi-compliant." If in step 300, a semi-compliant PCI device is not detected in any of the slots 36a-36c, execution of the routine terminates in step 302 and POST is completed in the normal fashion. Otherwise, execution proceeds to step 304, in which the entire address space is remapped so that it is forwarded to the secondary bus 22 via the bridge 12. In step 306, a

segment of VGA code stored in the memory of the host 14 is executed to display a message prompting the user to move the device from the secondary slot to a primary expansion slot. In step 308, a determination is made whether the device has been moved. If so, execution proceeds to step 302 and POST is completed; otherwise, execution remains at step 308 until the device is moved from the secondary slot. It will be recognized that, in operation, the solution illustrated in FIG. 3 will require a bypass such that, in case the detected device is in fact fully PCI compliant, the user can boot the PC normally without being forced to moving the device.

In a presently preferred alternative embodiment, the technique of the present invention is implemented using a hardware enhancement to the PCI-to-PCI bridge 12. As shown in FIG. 4, an enhancement circuit 400 for the PCI-to-PCI bridge 12 (FIG. 1) comprises a plurality of register sets 402, each of which includes a base register 402a and a limit register 402b. The base and limit registers 402a, 402b, of each set 402 may be programmed in a conventional manner to define an I/O address range. Although the enhancement circuit 400 shown in FIG. 4 comprises four register sets 402, it will be recognized that other numbers of register sets may be provided, depending on the particular embodiment of the PC 10.

Each I/O address received by the bridge 12 is stored in a next address register 404 and supplied to a comparator 406, which compares the next address from the register 404 with the I/O ranges defined by the register sets 402. If the next address falls within one of the defined I/O ranges, the next action taken will be determined by the contents of a configuration register 408, which register indicates, as two bits, a mode of operation for each of the register sets 402 individually. In particular, if the configuration register 408 contents indicate a first mode of operation (e.g., 01) for the register set 402 defining the I/O range in which the next address lies, the received I/O cycle is simply forwarded to the bus 22. Alternatively, if the configuration register contents correspond to a second mode of operation (e.g., 10) for the register set defining the I/O range in which the next address lies, an interrupt is generated on an interrupt pin (not shown) of the chip comprising the bridge 12 (FIG. 1), allowing the BIOS (FIG. 1, host 14) to handle an exception.

It will be recognized by those skilled in the art that the second mode of operation (i.e., the interrupt method) can be used to achieve an arbitrary granularity, as any I/O range can be mapped and software can handle selecting a finer granularity, independent of the operating system.

However, this method incurs a large performance hit, as it requires software intervention. In contrast, while the first mode of operation (i.e., the forwarding method) has a limited granularity, as a limited number of register sets 402 are used to define I/O ranges to be forwarded, it enables better and faster performance, as it is implemented completely in hardware and requires no software intervention.

It is understood that the present invention can take many forms and embodiments, the embodiments shown herein are intended to illustrate rather than limit, the invention, it being understood that variations may be made without departing from the spirit of the scope of the invention. For example, upon detection of a semi-compliant PCI device on the secondary PCI bus, the user could be instructed to move the device to the primary PCI bus in some manner other than via the display, such as illumination of an LED designated for that purpose or via an audio message. In addition, the present invention could be used in cases where the semi-compliant PCI device detected on the secondary bus, while not com-

pletely inoperable, would operate more efficiently on the primary bus, in which case the user could be informed, via a video or audio message or visual indicator, that the device would function better if moved to the primary bus.

Although illustrative embodiments of the invention have been shown and described, a wide range of modification, change and substitution is intended in the foregoing disclosure and in some instances some features of the present invention may be employed without a corresponding use of the other features. Accordingly, it is appropriate that the appended claims be construed broadly and in a manner consistent with the scope of the invention.

What is claimed is:

1. In a computer comprising a processor and a primary peripheral component interconnect ("PCI") bus electrically connected to a secondary PCI bus via a PCI-to-PCI bridge, a method comprising:

during a power on self test ("POST") of said computer, detecting a semi-compliant PCI device on said secondary PCI bus;

responsive to said detecting, causing said processor to trap on I/O writes to VGA address space, such that any attempted write cycle to said VGA address space will result in execution of an exception handler to notify a user to move said semi-compliant PCI device from said secondary PCI bus to said primary PCI bus;

suspending execution of said POST until said semi-compliant PCI device has been removed from said secondary PCI bus; and

completing said POST routine upon removal of said semi-compliant device from said secondary PCI bus.

2. The method of claim 1 wherein said notifying comprises displaying a video message on a display of said computer.

3. The method of claim 1 wherein said semi-compliant PCI device comprises a video card.

4. The method of claim 3 wherein said execution of said exception handler comprises:

activating an option ROM of said video card and reading an I/O address space thereof;

determining whether an address of said write cycle is part of an address space of said video card;

if said write cycle address is part of said address space of said video card, remapping said PCI-to-PCI bridge from an original mapping to handle said write cycle address and forwarding said write cycle to said video card.

5. The method of claim 4 wherein said execution of said exception handler further comprises generating an address look-up table, wherein said determining whether an address of said write cycle is part of an address space of said video card comprises determining whether said write cycle address is specified in said address look-up table.

6. The method of claim 4 wherein said execution of said exception handler further comprises:

initializing said video card; and

remapping said PCI-to-PCI bridge to said original mapping.

7. In a computer comprising a processor and a primary peripheral component interconnect ("PCI") bus electrically connected to a secondary PCI bus via a PCI-to-PCI bridge, an apparatus comprising:

means for detecting a semi-compliant PCI device on said secondary PCI bus during a power on self test ("POST") of said computer;

means responsive to said detecting, for causing said processor to trap on I/O writes to VGA address space, such that any attempted write cycle to said VGA address space will result in execution of an exception handler to notify a user to move said semi-compliant PCI device from said secondary PCI bus to said primary PCI bus;

means for suspending execution of said POST until said semi-compliant PCI device has been removed from said secondary PCI bus; and

means for competing said POST routine upon removal of said semi-compliant device from said secondary PCI bus.

8. The apparatus of claim 7 wherein said notifying means comprises a video message displayed on a display of said computer.

9. The apparatus of claim 7 wherein said semi-compliant PCI device comprises a video card.

10. The apparatus of claim 9 wherein said execution of said exception handler comprises:

activating an option ROM of said video card and reading an I/O address space thereof;

determining whether an address of said write cycle is part of an address space of said video card;

if said write cycle address is part of said address space of said video card, remapping said PCI-to-PCI bridge from an original mapping to handle said write cycle address and forwarding said write cycle to said video card.

11. The apparatus of claim 10 wherein said execution of said exception handler further comprises generating an address look-up table, wherein said determining whether an address of said write cycle is part of an address space of said video card comprises determining whether said write cycle address is specified in said address look-up table.

12. The apparatus of claim 10 wherein said execution of said exception handler further comprises:

initializing said video card; and

remapping said PCI-to-PCI bridge to said original mapping.

13. In a computer comprising a processor and a primary peripheral component interconnect ("PCI") bus electrically connected to a secondary PCI bus via a PCI-to-PCI bridge, a computer program stored on computer-readable media comprising:

instructions for detecting a semi-compliant PCI device on said secondary PCI bus during a power on self test ("POST") of said computer;

instructions responsive to said detecting for causing said processor to trap on I/O writes to VGA address space, such that any attempted write cycle to said VGA address space will result in execution of an exception handler to notify a user to move said semi-compliant PCI device from said secondary PCI bus to said primary PCI bus; and

instructions for allow said POST routine to be completed only upon removal of said semi-compliant device from said secondary bus.

14. The program of claim 13 wherein said instructions for notifying comprise instructions for displaying a video message on a display of said computer.

15. The program of claim 13 wherein said semi-compliant PCI device comprises a video card.

16. The program of claim 15 wherein said execution of said exception handler comprises:

activating an option ROM of said video card and reading an I/O address space thereof;

determining whether an address of said write cycle is part of an address space of said video card;

if said write cycle address is part of said address space of said video card, remapping said PCI-to-PCI bridge from an original mapping to handle said write cycle address and forwarding said write cycle to said video card.

17. The program of claim 15 wherein said execution of said exception handler further comprises generating an address look-up table, wherein said determining whether an address of said write cycle is part of an address space of said video card comprises determining whether said write cycle address is specified in said address look-up table.

18. The program of claim 15 wherein said execution of said exception handler further comprises:

initializing said video card; and

remapping said PCI-to-PCI bridge to said original mapping.

19. In a computer comprising a processor and a primary peripheral component interconnect ("PCI") bus electrically connected to a secondary PCI bus via a PCI-to-PCI bridge, a method comprising:

during a power on self test ("POST") of said computer, detecting a semi-compliant PCI device on said secondary PCI bus;

responsive to said detecting, suspending execution of said POST and remapping an address space of said computer such that a write cycle to any address in said address space is forwarded to said secondary PCI bus via said PCI-to-PCI bridge;

executing a segment of VGA code stored in a memory of said computer to display a message prompting a user to remove said semi-compliant PCI device from said secondary PCI bus; and

recommencing execution of said POST upon detection of removal of said semi-compliant PCI device from said secondary PCI bus.

20. The method of claim 19 wherein said semi-compliant PCI device is a video card.

21. In a computer comprising a processor and a primary peripheral component interconnect ("PCI") bus electrically connected to a secondary PCI bus via a PCI-to-PCI bridge, an apparatus comprising:

means for detecting a semi-compliant PCI device on said secondary PCI bus during a power on self test ("POST") of said computer;

means responsive to said detecting for suspending execution of said POST and remapping an address space of said computer such that a write cycle to any address in said address space is forwarded to said secondary PCI bus via said PCI-to-PCI bridge;

means for executing a segment of VGA code stored in a memory of said computer to display a message prompting a user to remove said semi-compliant PCI device from said secondary PCI bus; and

means for recommencing execution of said POST upon detection of removal of said semi-compliant PCI device from said secondary PCI bus.

22. The apparatus of claim 21 wherein said semi-compliant PCI device comprises a video card.

23. In a computer comprising a processor and a primary peripheral component interconnect ("PCI") bus electrically connected to a secondary PCI bus via a PCI-to-PCI bridge, a computer program stored on computer-readable medium comprising:

9

instructions for detecting a semi-compliant PCI device on said secondary PCI bus during a power on self test ("POST") of said computer;

instructions responsive to said detecting for suspending execution of said POST and remapping an address space of said computer such that a write cycle to any address in said address space is forwarded to said secondary PCI bus via said PCI-to-PCI bridge;

instructions for executing a segment of VGA code stored in a memory of said computer to display a message

10

prompting a user to remove said semi-compliant PCI device from said secondary PCI bus; and

instructions for recommencing execution of said POST upon detection of removal of said semi-compliant PCI device from said secondary PCI bus.

24. The program of claim 23 wherein said semi-compliant PCI device comprises a video card.

\* \* \* \* \*



UNITED STATES PATENT AND TRADEMARK OFFICE  
**CERTIFICATE OF CORRECTION**

PATENT NO. : 5,878,238  
DATED : March 2, 1999  
INVENTOR(S) : Doron Gan and Jeff Savage

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 7,

Line 58, please delete the word "allow" and insert therefor -- allowing --

Signed and Sealed this

Eighth Day of October, 2002

Attest:

A handwritten signature in black ink, appearing to read "James E. Rogan", with a horizontal line drawn underneath it.

Attesting Officer

JAMES E. ROGAN  
*Director of the United States Patent and Trademark Office*



US006311220B1

(12) **United States Patent**  
Fischer et al.

(10) Patent No.: **US 6,311,220 B1**  
(45) Date of Patent: **\*Oct. 30, 2001**

(54) **METHOD AND APPARATUS FOR  
COMMUNICATING DATA AND  
MANAGEMENT INFORMATION**

(75) Inventors: **Kelth S. Fischer**, Sunnyvale; **Corey A. Selby**, Hollister; **Robert W. Smith**, Los Altos, all of CA (US); **Kevin G. Smith**, Austin, TX (US)

(73) Assignee: **Nortel Networks Limited**, Montreal (CA)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **09/474,605**

(22) Filed: **Dec. 29, 1999**

#### Related U.S. Application Data

(63) Continuation of application No. 08/869,440, filed on Jun. 4, 1997, now Pat. No. 6,014,704.

(51) Int. Cl.<sup>7</sup> ..... **G06F 15/16**

(52) U.S. Cl. .... **709/230; 709/250**

(58) Field of Search ..... **709/250, 249,  
709/237, 230, 200**

(56) **References Cited**

#### U.S. PATENT DOCUMENTS

5,535,211	*	7/1996	Yano	370/85.2
5,550,802		8/1996	Worsley et al.	370/13
5,557,633	*	9/1996	Staab et al.	375/213
5,604,741		2/1997	Samueli et al.	370/402
5,727,170		3/1998	Mitchell et al.	710/105
5,754,552		5/1998	Allmond et al.	370/465
5,841,990		11/1998	Picazo, Jr. et al.	709/249
5,889,776	*	3/1999	Liang	370/389

\* cited by examiner


*Primary Examiner*—Mehmet B. Geckil


(74) *Attorney, Agent, or Firm*—Blakely Sokoloff Taylor & Zafman LLP

(57) **ABSTRACT**

A system is provided for communicating data between a first device and a second device across a communication link. Data is communicated between the first device and the second device across the communication link using a first protocol. Management information is communicated between the first device and the second device across the communication link using a second protocol. The system is capable of determining whether the first device is capable of communicating management information with the second device. If the first device is capable of communicating management information with the second device, then the system enables the communication of management information between the first device and the second device.

**23 Claims, 4 Drawing Sheets**



**FIG. 1**




FIG. 2

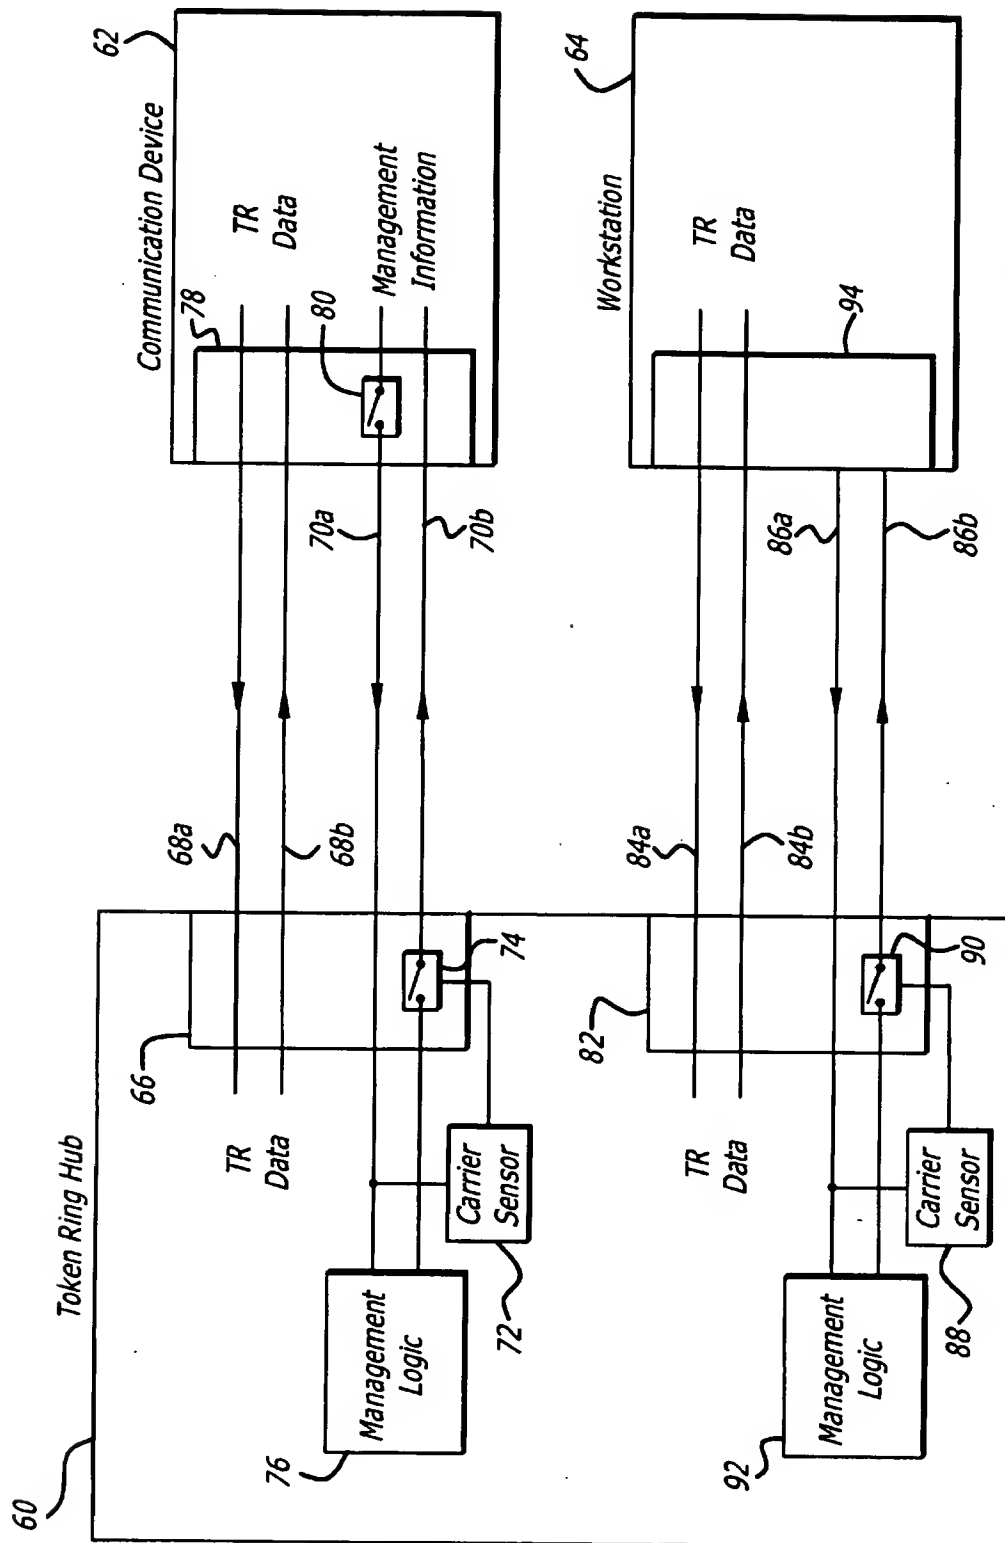
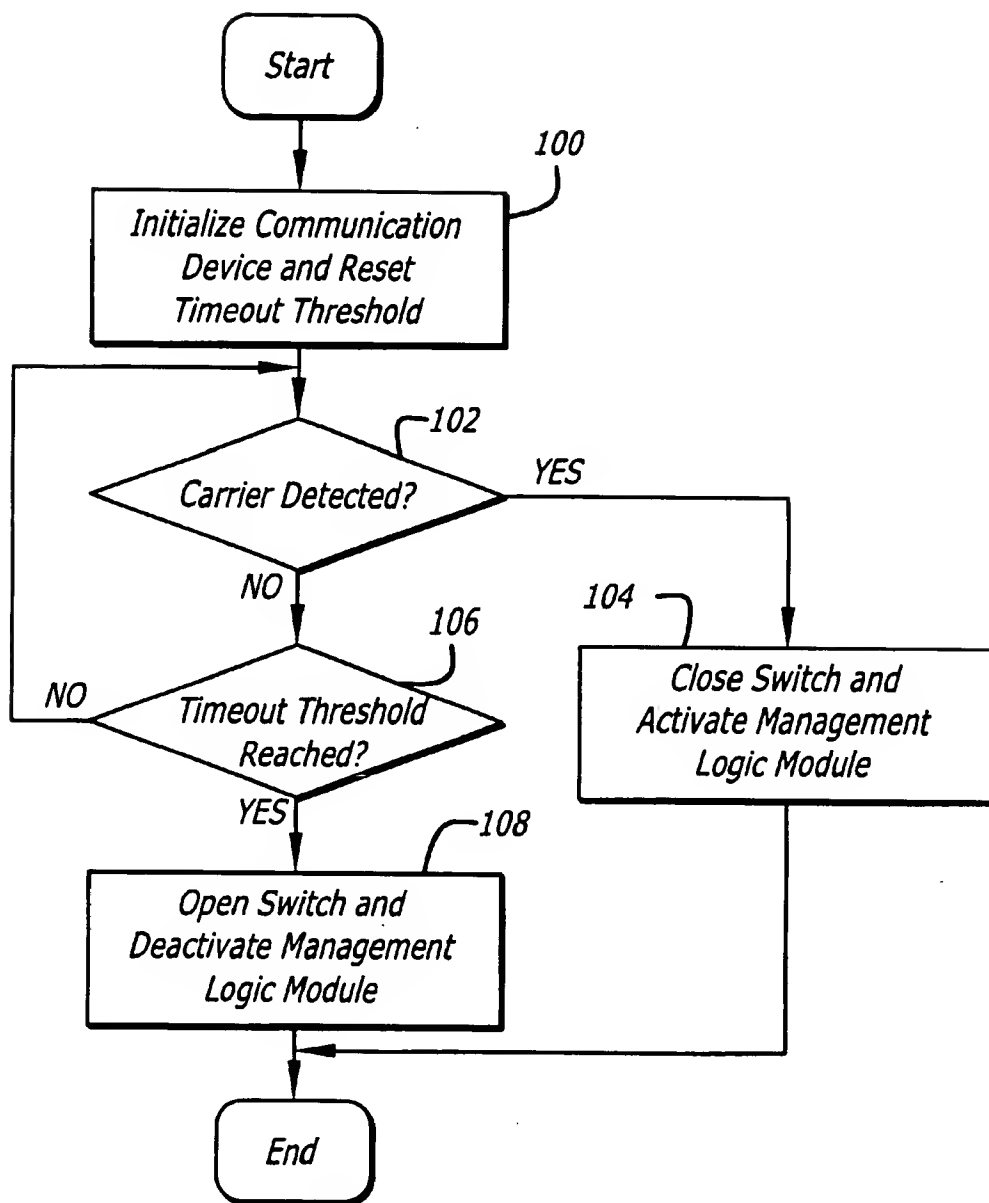


FIG. 3

**FIG. 4**

1

# METHOD AND APPARATUS FOR COMMUNICATING DATA AND MANAGEMENT INFORMATION

## CROSS-REFERENCES TO RELATED APPLICATIONS

This is a continuation of a U.S. patent application (application Ser. No. 08/869,440) filed Jun. 4, 1997 now U.S. Pat. No. 6,014,704.

## FIELD OF THE INVENTION

The present invention relates to systems for communicating data between two or more devices. More specifically, the invention provides a system for communicating data across a communication link using a first protocol and transmitting management information across the communication link using a second protocol.

## BACKGROUND

Various types of communication devices are available for communicating, managing, and handling data in a communication system such as a network. These communication devices include data terminal equipment (DTE), hubs, repeaters, network management units, and other devices capable of being coupled to another communication device or coupled to a network. Communication devices typically include one or more ports for transmitting and receiving data. These ports may be coupled to ports on other communication devices or coupled to a network.

In a network environment, multiple communication devices can be coupled together to permit the communication of data throughout the network. For example, a particular network may include several network hubs coupled to one another. In this example, each hub is capable of communicating with other hubs as well as communicating with and controlling communication devices coupled to the hub.

In known communication devices, a particular port may be designated as either a master port or a slave port (also referred to as a managed port or a controlled port). Network hubs typically have multiple ports for coupling to multiple communication devices. In certain situations, a network hub may be used as a control device for multiple communication devices. At the same time, the hub may receive control signals from another network device, such as a master device. To provide for this situation, known network hubs include multiple ports, one or more of which are dedicated to coupling a control device. Other ports in the network hub are dedicated to coupling communication devices that are controlled by the hub. Since all ports in the network hub have a dedicated configuration, a hub that is used only to control other communication devices cannot utilize the ports dedicated to coupling to a control device. Therefore, the hub resources are not fully utilized because the controlled port cannot be used to couple network devices.

Additionally, known systems use separate communication links for transmitting data and transmitting management or control signals. Thus, a first port is provided for communicating data with another communication device, and a second port is provided for communicating management or control information with the other communication device. This configuration requires two separate communication links and a pair of ports on each communication device. The additional communication links and ports increase the overall cost and complexity of the communication devices and

2

require additional physical space on the device to attach the additional ports.

It is therefore desirable to provide a communication system that is capable of transmitting both data and management information across a single communication link using a single communication port on each communication device.

## SUMMARY OF THE INVENTION

Embodiments of the present invention provide a communication system that is capable of transmitting both data and management information across a single communication link using a single communication port on each communication device. The communication port determines the type of port at the opposite end of the communication link and configures its operation accordingly. Two different protocols can be used to communicate the data and the management information across the communication device.

A particular embodiment of the invention communicates data between a first device and a second device across a communication link using a first protocol. Management information is communicated between the first device and the second device across the communication link using a second protocol.

In other embodiments of the invention, a communication port is determined whether the first device is capable of communicating management information with the second device. If the first device is capable of communicating management information with the second device, then the communication port is enabled to communicate management information between the first device and the second device using the second protocol.

## BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example in the following drawings in which like references indicate similar elements. The following drawings disclose various embodiments of the present invention for purposes of illustration only and are not intended to limit the scope of the invention.

FIG. 1 illustrates an embodiment of a communication device having a dual-mode port and a fixed-mode port.

FIG. 2 illustrates an example arrangement of multiple devices coupled to a communication device.

FIG. 3 illustrates an embodiment of a Token Ring hub coupled to multiple Token Ring devices.

FIG. 4 is a flow diagram illustrating an embodiment of a procedure for configuring a dual-mode port.

## DETAILED DESCRIPTION

The following detailed description sets forth numerous specific details to provide a thorough understanding of the invention. However, those of ordinary skill in the art will appreciate that the invention may be practiced without these specific details. In other instances, well-known methods, procedures, protocols, components, and circuits have not been described in detail so as not to obscure the invention.

Embodiments of the present invention provide a communication system that is capable of transmitting both data and management information across a single communication link using a single communication port on each communication device. The communication port automatically senses the type of port at the opposite end of the communication link and configures its operation accordingly. The commu-

nication port is configured to communicate only data if the port on the opposite end of the communication link is incapable of communicating management information. Additionally, the communication port is configured to communicate data as well as management information if the port on the opposite end of the communication link is capable of communicating management information. In this configuration, one protocol can be used to communicate the data and another protocol can be used to communicate the management information. In certain situations, there may be economic or technical advantages in using existing hardware and software associated with one protocol (e.g., Ethernet) to manage the flow of data, which uses a different protocol (e.g., Token Ring).

FIG. 1 illustrates an embodiment of a communication device 10 having a dual-mode port 12 and a fixed-mode port 14. Although communication device 10 is shown with two ports, it will be appreciated that communication device 10 may include any number of dual-mode ports 12 and any number of fixed-mode ports 14. A particular embodiment of communication device 10 includes six dual-mode ports 12 and eighteen fixed-mode ports 14, for a total of twenty-four ports. Other embodiments of the invention may include one or more dual-mode ports 12, but not include any fixed-mode ports.

Fixed-mode port 14 has a single mode of operation. This single mode of operation may include communicating data using a particular protocol, such as Ethernet, Token Ring, or any other data communication protocol. In a particular embodiment of the invention, fixed-mode port 14 is a data communication port using the Token Ring protocol. Fixed-mode port 14 communicates data and other information using a communication link 20.

Dual-mode port 12 has multiple modes of operation. For example, in a first mode of operation, dual-mode port 12 is a data communication port using the Token Ring protocol. In this first mode of operation, port 12 operates as a conventional Token Ring port. In a second mode of operation, dual-mode port 12 communicates data using the Token Ring protocol and also communicates management information using another protocol, such as Ethernet. As shown in FIG. 1, in this dual mode of operation, port 12 communicates data using a communication link 18 and communicates management information using a communication link 16. The management information communicated across link 16 is generated by management logic 22, which is coupled to dual-mode port 12. This management information may include commands and other information necessary to control or monitor the operation of other communication devices coupled to communication device 10. In an embodiment of the invention, the management information includes statistics collected (e.g., number of packets) and the commands include enabling or disabling other ports in the communication device.

Although communication links 16 and 18 are illustrated in FIG. 1 as two separate communication links, they may be incorporated into a single cable or other transmission medium. For example, communication link 16 may use a pair of conductors in a twisted-pair cable, and communication link 18 may use another pair of conductors in the same twisted-pair cable. Thus, instead of using separate cables (and separate ports) to transmit data and management information, the embodiment of FIG. 1 permits a single cable (and a single port) to transmit both data and management information. Additionally, different protocols may be used with each pair of conductors (e.g., the Token Ring protocol can be used to transmit data, and the 10 Base T

Ethernet protocol can be used to transmit management information). Since the management information is transmitted on a separate pair of conductors, the information does not interfere with or affect the bandwidth of the data being transmitted on the other conductors, if properly designed. In a particular embodiment of the invention, communication links 16 and 18 are contained in a single cable having multiple twisted pairs of conductors (e.g., category 5 cable).

A data transmit and receive module 24 is coupled to both dual-mode port 12 and fixed-mode port 14. Module 24 handles the incoming and outgoing data communicated across communication links 18 and 20 by ports 12 and 14, respectively.

Although not shown in FIG. 1, communication device 10 may also include additional components or modules, such as a Central Processing Unit (CPU), memory devices, mass storage devices, additional ports, and other modules.

FIG. 2 illustrates an example arrangement of multiple devices coupled to a communication device 30. Various modules and components of communication device 30 have been omitted for purposes of clarity. Communication device 30 is coupled to multiple devices 32, 34, and 36. Communication device 30 includes a dual-mode port 38 and a pair of fixed-mode Token Ring ports 40 and 42. Dual-mode port 38 is capable of operating as a fixed-mode Token Ring port for communicating Token Ring data, or operating as a dual-mode port for communicating Token Ring data as well as management information, as discussed above. In the embodiment of FIG. 2, dual-mode port 38 is configured to transmit Token Ring data across a communication link 46 and is configured to transmit management information across a communication link 48 using the Ethernet protocol. As discussed above with respect to FIG. 1, communication links 46 and 48 may be contained within a single cable, such as a twisted-pair cable. Although a particular example is described as using the Token Ring protocol to transmit data and the Ethernet protocol to transmit management information, those of ordinary skill in the art will appreciate that other protocols may be used to transmit data or management information.

Communication device 32 includes a dual-mode port 44 that is capable of communicating data (such as Token Ring data) using communication link 46 and capable of communicating management information (e.g., using the Ethernet protocol) on communication link 48. Communication device 32 may be a hub (e.g., a Token Ring hub) or any other device capable of communicating with communication device 30. Although not shown in FIG. 2, other devices (such as a workstation) may be coupled to communication device 32.

Workstation 34 includes a fixed-mode Token Ring port 50 coupled to communication link 52. Thus, workstation 34 is capable of communicating Token Ring data with communication device 30. However, since workstation 34 contains a fixed-mode port, communication link 52 does not support the communication of management information between communication device 30 and workstation 34. Workstation 34 may be any type of workstation using any type of processor. Similarly, workstation 36 includes a fixed-mode Token Ring port 54 coupled to communication link 56. As with workstation 34 discussed above, workstation 36 is capable of communicating Token Ring data with communication device 30, but does not support the communication of management information with device 30. Workstations 34 and 36 may be any type of station coupled to a Token Ring network, such as a server or a network printer.

FIG. 3 illustrates an embodiment of a Token Ring hub 60 coupled to multiple Token Ring devices 62 and 64. Token



5

Ring hub 60 includes two dual-mode ports 66 and 82 coupled to communication device 62 and workstation 64, respectively. Dual-mode port 66 is coupled to four conductors 68a, 68b, 70a, and 70b. In an embodiment of the invention, conductors 68a, 68b, 70a, and 70b are grouped together in a single cable. Conductors 68a and 68b are used to transmit Token Ring data (labeled "TR Data") between Token Ring hub 60 and communication device 62. Conductors 70a and 70b are used to transmit management information between Token Ring hub 60 and communication device 62. As discussed above, an embodiment of the invention uses the Ethernet protocol to transmit management information between Token Ring hub 60 and communication device 62.

In an alternative embodiment of the invention, each conductor 68a-68d in FIG. 3 is replaced with a twisted pair of conductors, resulting in a total of eight conductors. All four twisted pairs may be included in a single cable.

Token Ring hub 60 includes a carrier sensor 72 coupled to conductor 70a and a switch 74. Carrier sensor 72 determines whether a carrier (e.g., a link test pulse in an Ethernet-based system) is present on conductor 70a, which communicates management information. The carrier signal type will vary with the protocol used to communicate across conductors 70a and 70b. This carrier is part of the management information and is separate from any carrier transmitted on a conductor which communicates data. If a carrier is detected, then carrier sensor 72 closes switch 74 to permit communication of management information on conductors 70a and 70b. If a carrier is not detected within a particular time period, then carrier sensor 72 opens switch 74 to prevent transmission of management information across conductor 70b. This detection of a carrier on conductor 70a is performed automatically by carrier sensor 74, thereby providing automatic configuration of dual-mode port 66. Switch 72 can be implemented using one or more logic gates, a relay, a transistor, or other switching device.

If communication device 62 includes a dual-mode port (as shown in FIG. 3) capable of communicating management information across conductors 70a and 70b, then carrier sensor 72 will detect the presence of a carrier on conductor 70a and configure port 66 as a dual-mode port. In this situation, a management logic module 76 controls the flow of management information between Token Ring hub 60 and communication device 62. If communication device 62 does not contain a dual-mode port, or if the port is configured to transmit data only, then carrier sensor 72 will not detect a carrier on conductor 70a and will configure port 66 as a data-only port. Thus, various types of devices having different types of ports can be coupled to port 66 without requiring any manual configuration of port 66.

Communication device 62 includes a dual-mode port 78 capable of operating as a data-only port or a dual-mode port for transmitting data as well as management information. The mode of operation of port 78 is determined by a switch 80. In a particular embodiment of the invention, switch 80 is a manual switch that is set by a user of device 62 or a system administrator, indicating the mode in which port 78 will operate. Alternatively, switch 80 may be any type of manual, automatic, or software-controlled switch. If switch 80 is open, then port 78 operates as a fixed-mode port because there is a physical break in the connection to conductor 70a such that no carrier will be transmitted to Token Ring hub 60. Alternatively, port 78 may operate as a fixed-mode port if the carrier on conductor 70a is suppressed. If switch 80 is closed, then port 78 operates as a dual-mode port, thereby permitting the simultaneous

6

exchange of Token Ring data and management information between Token Ring hub 60 and communication device 62.

Token Ring hub 60 also includes an additional dual-mode port 82 similar to port 66 discussed above. Dual-mode port 82 communicates Token Ring data with workstation 64 using a pair of conductors 84a and 84b. Dual-mode port 82 is also coupled to conductors 86a and 86b, which are used to communicate management information between Token Ring hub 60 and workstation 64. In one embodiment of the invention, conductors 84a, 84b, 86a, and 86b are grouped together in a single twisted-pair cable.

A carrier sensor 88 in Token Ring hub 60 is coupled to conductor 86a and a switch 90. As discussed above with reference to carrier sensor 72, if a carrier is detected on conductor 86a, carrier sensor 88 closes switch 90 to permit the communication of management information across conductors 86a and 86b. Otherwise, carrier sensor 88 opens switch 90 to prevent the transmission of management information across conductor 86b. A management logic module 92 controls the flow of management information between Token Ring hub 60 and workstation 64.

In the example shown in FIG. 3, workstation 64 contains a fixed-mode port 94 that is not capable of receiving management information from Token Ring hub 60. In this example, workstation 64 will not transmit a carrier across conductor 86a. Therefore, carrier sensor 88 will open switch 90 to prevent the transmission of management information across conductor 86b. Preventing the unnecessary transmission of management information across conductor 86b reduces electromagnetic interference.

Thus, as shown in FIG. 3, dual mode ports 66 and 82 are automatically configured to communicate with the device coupled to the port. If the device coupled to the port (e.g., communication device 62) is capable of communicating management information, then the port is configured to support this mode of operation. However, if the device coupled to the port (e.g., workstation 64) is not capable of communicating management information, then the port is configured to support a data-only mode of operation. This automatic configuration permits a single port to be coupled to various types of devices and ports without requiring a manual configuration of the port or requiring that each port have a dedicated mode of operation.

FIG. 4 is a flow diagram illustrating an embodiment of a procedure for configuring a dual-mode port, such as port 66 or 82 in FIG. 3. At block 100, the communication device is initialized and a timeout threshold is reset. The timeout threshold, as discussed below, identifies the length of time that the dual-mode port will search for a carrier (indicating that the port on the opposite end of the communication link is capable of communicating management information). At block 102, the procedure determines whether a carrier has been detected (e.g., on conductor 70a or 86a in FIG. 3). If a carrier is detected, then block 104 closes the switch in the dual-mode port and activates the management logic module. At this point, the dual-mode port is configured to communicate data (e.g., Token Ring data) as well as management information across the communication link.

If a carrier has not been detected at block 102, then the procedure continues to block 106 to determine whether the timeout threshold has been reached. The timeout threshold identifies the length of time that the dual-mode port will search for a carrier. If the threshold has not been reached in block 106, the procedure returns to step 102 to continue searching for a carrier. If the timeout threshold is reached or exceeded at block 106, then the procedure continues to step 108 where the switch in the dual-mode port is opened and the management logic module is deactivated.

The procedure illustrated in FIG. 4 can be executed periodically to check the status of the port at the opposite

end of a communication link. Thus, if a device is disconnected from the communication link, or a different device is connected to the communication link, the procedure of FIG. 4 identifies this change and modifies the switch setting in the dual-mode port, if necessary. In a particular embodiment of the invention, the procedure of FIG. 4 is executed approximately once each second for each dual mode port in the communication device.

In alternative embodiments, the present invention may be applicable to implementations of the invention in integrated circuits or chip sets, wireless implementations, switching systems products and transmission systems products. For purposes of this application, the term switching systems products shall be taken to mean private branch exchanges (PBXs), central office switching systems that interconnect subscribers, toll/tandem switching systems for interconnecting trunks between switching centers, and broadband core switches found at the center of a service provider's network that may be fed by broadband edge switches or access muxes, and associated signaling, and support systems and services. The term transmission systems products shall be taken to mean products used by service providers to provide interconnection between their subscribers and their networks such as loop systems, and which provide multiplexing, aggregation and transport between a service provider's switching systems across the wide area, and associated signaling and support systems and services.

From the above description and drawings, it will be understood by those of ordinary skill in the art that the particular embodiments shown and described are for purposes of illustration only and are not intended to limit the scope of the invention. Those of ordinary skill in the art will recognize that the invention may be embodied in other specific forms without departing from its spirit or essential characteristics. References to details of particular embodiments are not intended to limit the scope of the claims.

What is claimed is:

1. A method comprising:

coupling a first communication link and a second communication link to a first port of a first device; and configuring the first port to transfer data over the first communication link and management information over the second communication link in response to a carrier signal being detected on the second communication link.

2. The method of claim 1, wherein the first communication link and the second communication link are included in one communication cable.

3. The method of claim 1, wherein the management information includes commands to control operations of a second device coupled to the second communication link.

4. The method of claim 1, wherein the management information includes information to monitor a second device coupled to the second communication link.

5. The method of claim 1, wherein configuring the first port includes configuring the first port to transfer data in accordance with a first protocol and to transfer the management information in accordance with a second protocol.

6. The method of claim 5, wherein the transfer of the data and the transfer of the management information are generally concurrent to each other.

7. The method of claim 5, wherein the first protocol is Token Ring.

8. The method of claim 7, wherein the second protocol is Ethernet.

9. The method of claim 1 further comprising:

configuring the first port to only transfer data over the first communication link if the carrier signal is not detected on the second communication link.

10. A communication device comprising:

a port means for transferring information between said communication device and another device; and

configuration means for configuring the port means to transfer both data and management information, if it is determined that said another device is capable of communicating management information.

11. The communication device of claim 10 further comprising means for coupling the port means to said another device, the means for coupling includes at least one communication cable having a first communication link and a second communication link.

12. The communication device of claim 11, wherein the at least one communication cable is a twisted-pair cable including the first communication link having a first pair of conductors and the second communication link having a second pair of conductors.

13. The communication device of claim 11, wherein the configuration means for configuring the port includes:

sensing means for automatically detecting whether a carrier signal is present on the second communication link, the carrier signal indicating that said another device is capable of communicating management information; and

switching means for configuring the port means to transfer both the data and management information in response to the carrier signal being detected.

14. The communication device of claim 10, wherein the management information includes commands to control operations of said another device.

15. The communication device of claim 10, wherein the management information includes information to monitor said another device.

16. A communication device comprising:

a port to communicate data over a first communication link and management information over a second communication link;

a carrier sensor coupled to the port, the carrier sensor to detect a carrier signal provided through the port; and

a switch coupled to the port and the carrier sensor, the switch to permit communication of both the data and the management information when the carrier signal is detected.

17. The communication device of claim 16, wherein the switch also prohibits communication of the management information when the carrier signal is not detected.

18. The communication device of claim 16 wherein the port communicates data according to a first protocol and management information according to a second protocol.

19. The communication device of claim 18, wherein the first communication link and the second communication link are included in a single cable having a plurality of conductors.

20. The communication device of claim 18, wherein the first communication link and the second communication link are included in a twisted-pair cable having a plurality of pairs of conductors.

21. The communication device of claim 20, wherein the port communicates the data in accordance with the first protocol across a first pair of conductors and the port communicates the management information in accordance with the second protocol across a second pair of conductors.

22. The communication device of claim 16 being a switching system product.

23. The communication device of claim 16 being a transmission system product.

\* \* \* \* \*